



NetServ: Deploying Customized Network Services on Demand

Henning Schulzrinne,
Jae Woo Lee & Suman Srinivasan
Columbia University

Joint work with: Bell Labs (Alcatel-Lucent), Deutsche Telekom,
DOCOMO Euro-Labs



KuVS Fachgespräch Fraunhofer FOKUS



NetServ overview

Extensible architecture for core network services




Figure 3: Instantiation of services over tunable building blocks.

Modularization

- Building Blocks
- Service Modules

Virtual services framework

- Security
- Portability

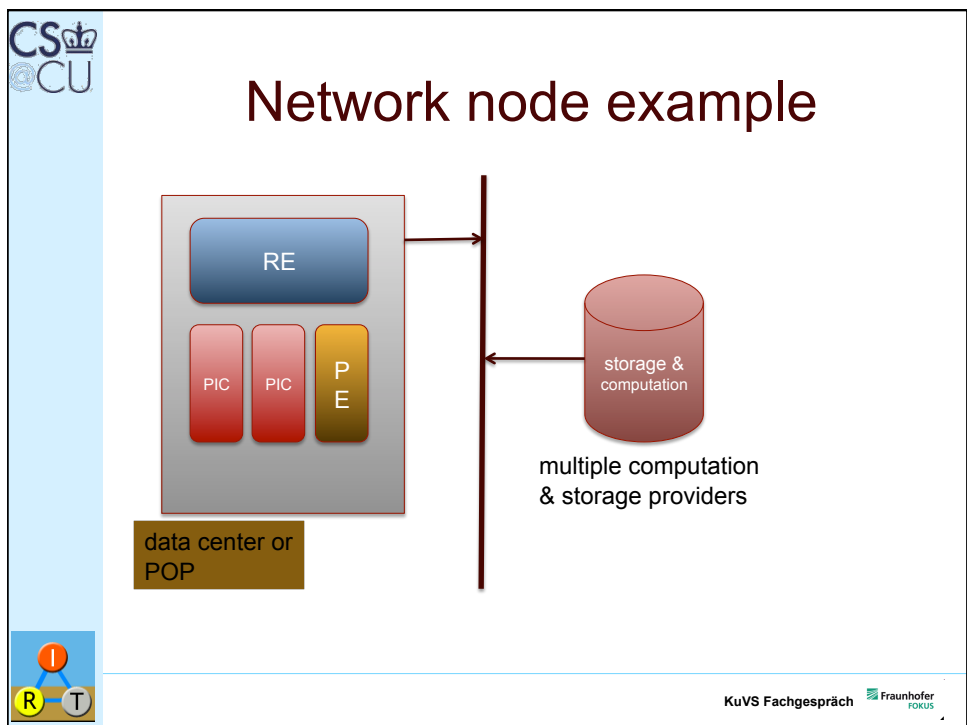
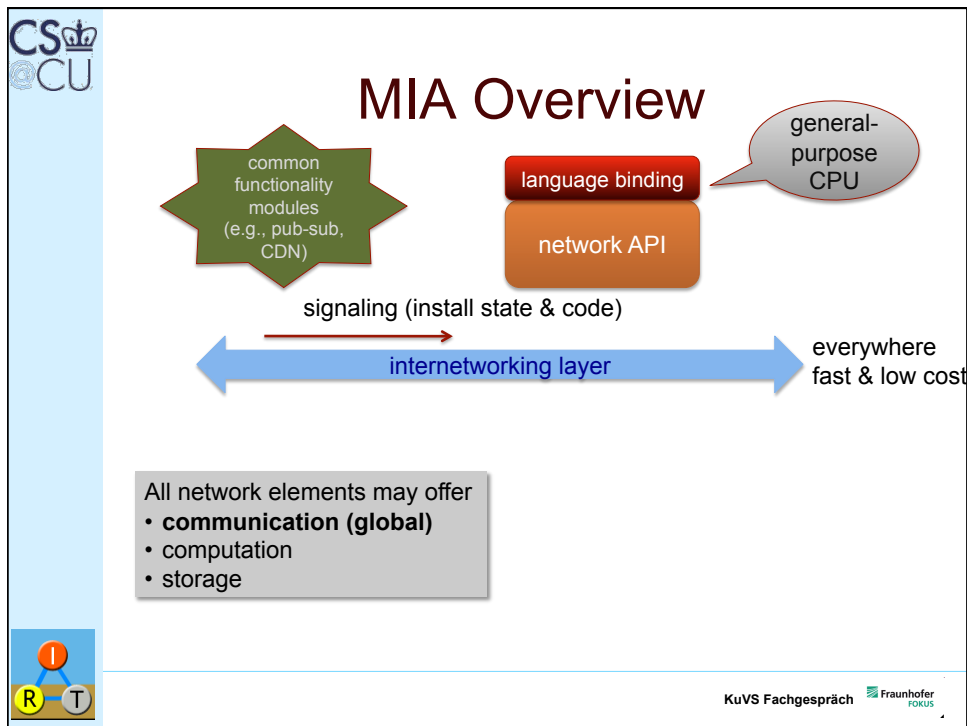
NSF FIND four-year project

- Columbia University
- Bell Labs
- Deutsche Telekom
- DOCOMO Euro-Labs

No more *ossification* in NGI



KuVS Fachgespräch Fraunhofer FOKUS





Different from Active Networks?

- Active Networks
 - Packet contains executable code
 - Can modify router states and behavior
 - Not successful
 - Per-packet processing too expensive
 - Security concerns
 - No compelling *killer app* to warrant such a big shift
 - Notable work: ANTS, Janos, Switchware
- NetServ
 - Virtualized services on current, passive networks
 - Service invocation is signaling driven, not packet driven
 - Service modules are stand-alone, addressable entities
 - Separate from packet forwarding plane
 - Extensible plug-in architecture



Building Blocks

- Key components of network services
 - Access to network-level resource
 - Implementation of common functionality
- For example:
 - Link monitoring and measurement
 - Routing table
 - Packet capture
 - Data storage and lookup





Service Modules

- Full-fledged service implementations
 - Use Building Blocks and other Service Modules
 - Can be implemented across multiple nodes
 - Invoked by applications
- Examples:
 - Routing-related services
 - Multicast, anycast, QoS-based routing
 - Monitoring services
 - Link & system status, network topology
 - Identity services
 - Naming, security
 - Traffic engineering services
 - CDN, redundancy elimination, p2p network support



Deployment Scenarios

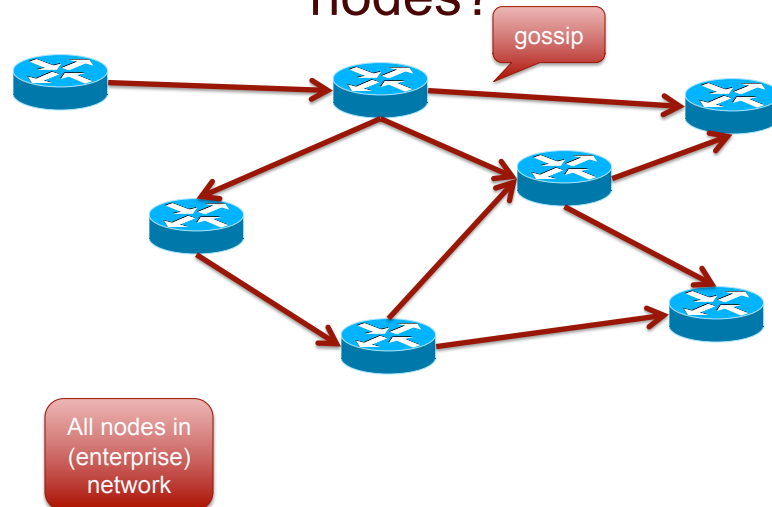
- Three actors
 - Content publisher (e.g. youtube.com)
 - Service provider (e.g. ISP)
 - End user
- Model 1: Publisher-initiated deployment
 - Publisher rents router space from providers (or end users)
- Model 2: Provider-initiated deployment
 - Publisher writes NetServ module
 - Provider sees lots of traffic, fetches and installs module
 - Predetermined module location (similar to robots.txt)
- Model 3: User-initiated deployment
 - User installs NetServ module to own home router or PC

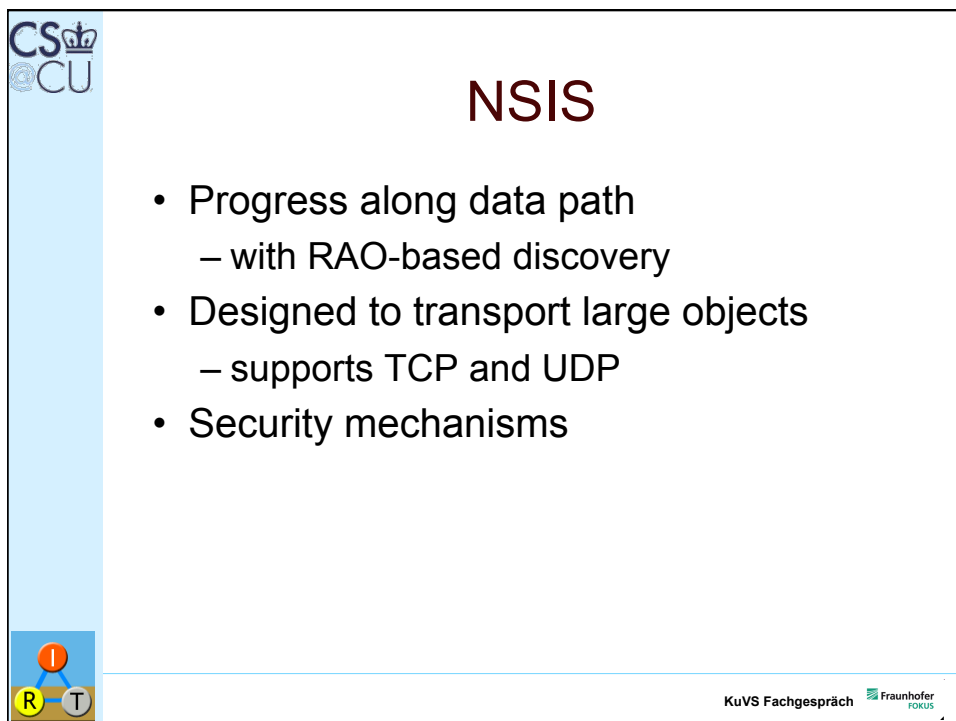
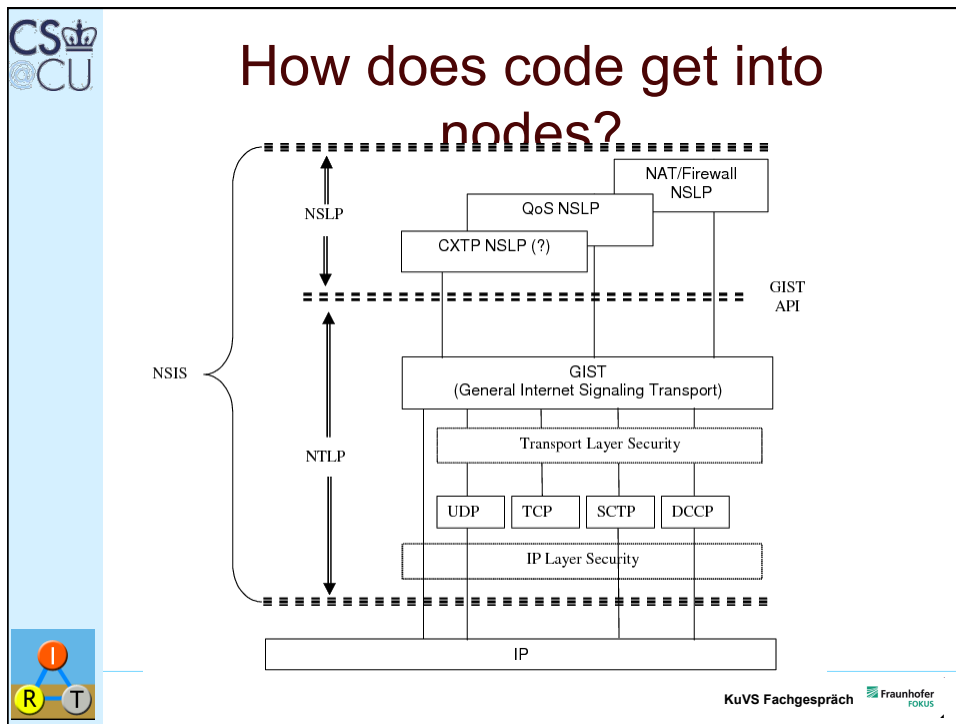


Where does code run?

- All (or some?) nodes in a network
 - AS, enterprise LAN
- Some or all nodes along path
 - data path from source to destination
- Selected nodes by property
 - e.g., one in each AS

How does code get into nodes?







First prototype implementation

- Proof-of-concept for dynamic network service deployment
 - Open-source Click modular router
 - Java OSGi dynamic module system
- Promising initial measurement results
 - NetServ overhead acceptable compared to other overhead



Technology: Click router

- Runs as a Linux kernel module or user-level program
- Modules written in C++ (called *Elements*) are configured in a text file
- Elements are arranged in a directed graph, through which packets traverse
- Example:
 - Click router command:


```
sudo click print.click
```
 - Configuration file print.click:


```
FromDevice(en0)->CheckIPHeader(14)->IPPrint->Discard;
```
- <http://www.read.cs.ucla.edu/click/>

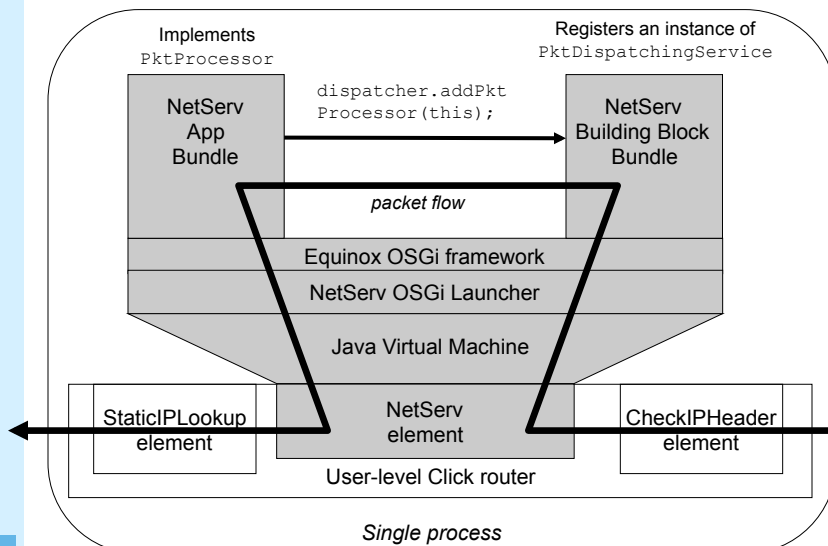


Technology: OSGi

- Dynamic module system for Java
 - Modules loaded and unloaded at runtime
 - *Bundle*: self-contained JAR file with specific structure
 - Open-source implementations: Apache Felix, Eclipse Equinox
- Security and accounting
 - Security built on Java 2 Security model
 - Permission-based access control
 - No fine-grained control or accounting for CPU, storage, bandwidth
 - Can load native code with appropriate permission
 - Strict separation of bundles
 - Classpath set up by Bundle class loader
 - Inter-bundle communication only through published interfaces



1st prototype implementation



CS@CU

Demo: NetServ prototype

(1) Regular Incoming packets

(2) "Operator" can view modules on router

(3) Operator loads a new module (that makes all data uppercase)

(4) Packets are modified

(5) Operator stops the module

(6) No more packet modification

Listening on port 7777 ...
 Aug 14, 2009 11:44:44 AM NetServ.Builder:BuildingBlock.Activator:Start
 INFO: registering PktDispatchingService with (PktDispatchingService) [1250264684.542293: length 36 135.112.131.10.369731 > 135.112.130.8.44444: udp 16 netServ.
 1250264684.508154: length 36 135.112.131.10.60372 > 135.112.130.8.44444: udp 16 netServ.
 1250264694.528855: length 36 135.112.131.10.42839 > 135.112.130.8.44444: udp 16 netServ.
 1250264699.548799: length 36 135.112.131.10.98024 > 135.112.130.8.44444: udp 16 netServ.
 Aug 14, 2009 11:45:02 AM NetServ.Builder:BuildingBlock.Activator:Start
 INFO: registering PktDispatchingService with (PktDispatchingService) [1250264709.569479: length 36 135.112.131.10.36980 > 135.112.130.8.44444: udp 16 netServ.
 1250264709.569479: length 36 135.112.131.10.36980 > 135.112.130.8.44444: udp 16 netServ.
 Aug 14, 2009 11:45:11 AM NetServ.Builder:BuildingBlock.Activator:Start
 INFO: registering PktDispatchingService with (PktDispatchingService) [1250264714.610106: length 36 135.112.131.10.36980 > 135.112.130.8.44444: udp 16 netServ.
 1250264714.610106: length 36 135.112.131.10.36980 > 135.112.130.8.44444: udp 16 netServ.

```

Framework is launched.
id      State  Bundle
0      ACTIVE org.eclipse.osgi_3.5.0.v20090311-1300
1      ACTIVE NetServ.Builder:BuildingBlock_0.0.1

omg| install file:///home/sumans/netServ/prot0-02/example1.jar
Bundle id is 4

omg| ss

Framework is launched.
id      State  Bundle
0      ACTIVE org.eclipse.osgi_3.5.0.v20090311-1300
1      ACTIVE NetServ.Builder:BuildingBlock_0.0.1
4      INSTALLED example.app1.UDPProcessor_0.0.1

omg| start 4

omg| stop 4

omg|
  
```

KuVS Fachgespräch Fraunhofer FOKUS

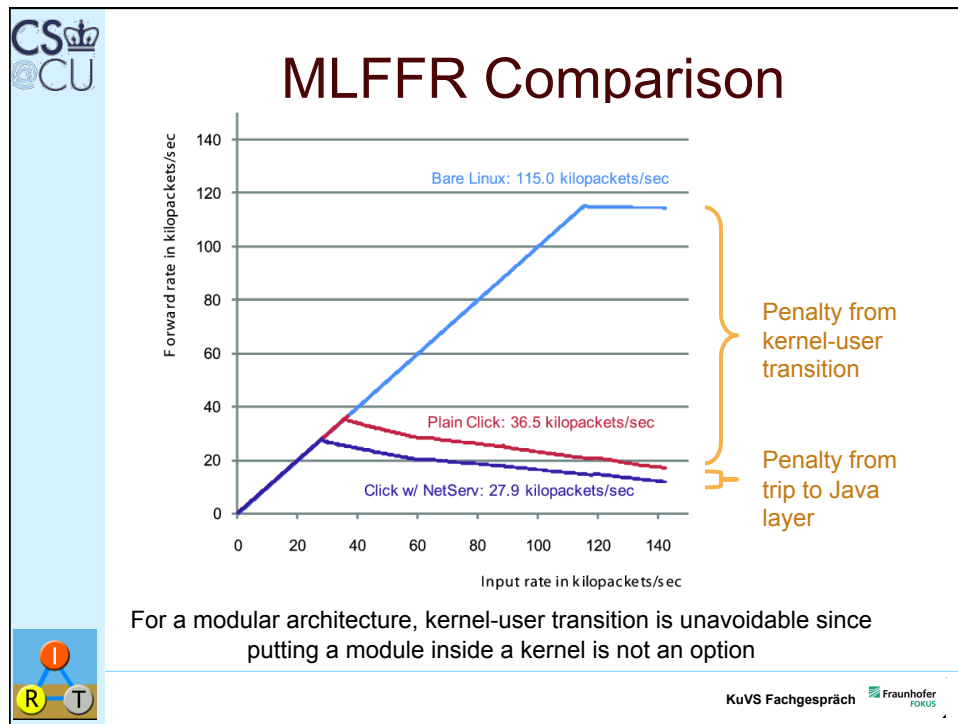
CS@CU

Performance Evaluation

- Initial measurements on the first prototype
 - NetServ on user-level Click router
 - Maximum Loss Free Forward Rate (MLFFR)
- Future work on next-generation prototypes
 - NetServ on JUNOS, kernel-mode Click
 - Ping latency
 - Microbenchmarks
 - Throughput for non-trivial services

18

KuVS Fachgespräch Fraunhofer FOKUS



Current Work: CDN on NetServ

- On-Path CDN
 - Prototype implemented
- Dynamic content migration
 - Moving content closer to the end user according to demand
- Building blocks
 - Network monitoring
 - Content discovery
 - Caching proxy

KuVS Fachgespräch Fraunhofer FOKUS



API examples

- Avoid SNMP retrieval problems
 - all or nothing (typical)
 - hard to do selective triggers
- Flow management
 - counters, measurement
- System information
 - like system MIB: geo location, uptime, interface speeds, ...
 - routing table
 - routing table changes (“tell me if route to X changes”)



Current Work: NetServ Platform

- Ubiquitous NetServ
 - From big to small devices
 - Real router: Juniper’s JUNOS
 - Personal computer: Kernel-mode Click
 - Home router: Linux using iptables
- Security and resource control
 - Enable various deployment scenarios
 - Support different economic incentives





Related Work

- Cisco's Programmable Overlay Router
- Juniper's JUNOS SDK
- DaVinci project
- VROOM (virtual routers on the move)
- OpenFlow Switch
- Ethane



Future Internet Architecture?

- Really closer to urban design
 - zoning, fire codes and infrastructure (rail, water)
 - plus oversight (fire marshal & building inspector)
 - architecture changes, urban designs stay
 - see Washington, DC & Berlin
- “Architecture” must be
 - expressible in one sentence
 - avoid limiting options (unknown unknowns)
 - avoid imposing unnecessary costs





The network services fallacy

- We tried adding network services as protocols:
 - multicast
 - QoS
 - mobility
 - security
- All were, more-or-less, failures
 - (or underperformed expectations)
 - hard to secure, not quite right



Thoughts on architecture

- Long-term constant: service model
 - equivalent of railroad track & road width
- Identify core functions we need
 - routing
 - congestion control
 - name lookup
 - path state establishment
 - ...
- Learn from history
 - why didn't these get done "right"?
- Need engineering principles
- Requirement list doesn't help



